



US009086924B2

(12) **United States Patent**  
**Barsness et al.**

(10) **Patent No.:** **US 9,086,924 B2**  
(45) **Date of Patent:** **Jul. 21, 2015**

(54) **EXECUTING A DISTRIBUTED JAVA APPLICATION ON A PLURALITY OF COMPUTE NODES**

(75) Inventors: **Eric L. Barsness**, Rochester, MN (US);  
**David L. Darrington**, Rochester, MN (US); **Amanda E. Peters**, Rochester, MN (US); **John M. Santosuosso**, Rochester, MN (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 2213 days.

(21) Appl. No.: **12/109,238**

(22) Filed: **Apr. 24, 2008**

(65) **Prior Publication Data**  
US 2009/0271784 A1 Oct. 29, 2009

(51) **Int. Cl.**  
**G06F 9/455** (2006.01)  
**G06F 9/50** (2006.01)  
**G06F 9/48** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 9/5044** (2013.01); **G06F 9/45558** (2013.01); **G06F 9/4843** (2013.01); **G06F 2209/508** (2013.01); **G06F 2209/5017** (2013.01)

(58) **Field of Classification Search**  
None  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,842,759 B2 \* 1/2005 Haggar et al. .... 1/1  
7,703,089 B2 \* 4/2010 Birenheide .... 717/166

8,296,759 B1 \* 10/2012 Hutchins et al. .... 718/1  
2003/0120776 A1 \* 6/2003 Avvari et al. .... 709/225  
2004/0154016 A1 \* 8/2004 Randall .... 718/1  
2005/0005200 A1 \* 1/2005 Matena et al. .... 714/38  
2006/0085785 A1 \* 4/2006 Garrett .... 718/1  
2006/0271395 A1 \* 11/2006 Harris et al. .... 705/1  
2007/0130566 A1 \* 6/2007 van Rietschote et al. .... 718/1  
2008/0028401 A1 \* 1/2008 Geisinger .... 718/1  
2008/0034365 A1 \* 2/2008 Dahlstedt .... 718/1

#### OTHER PUBLICATIONS

Wenzhang Zhu, Cho-Li Wang, and Francis C. M. Lau. "Lightweight Transparent Java Thread Migration for Distributed JVM" International Conference on Parallel Processing, 2003, pp. 465-472.\*

\* cited by examiner

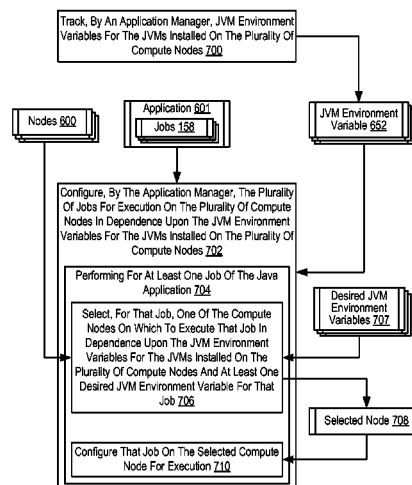
Primary Examiner — Jacob A Petranek

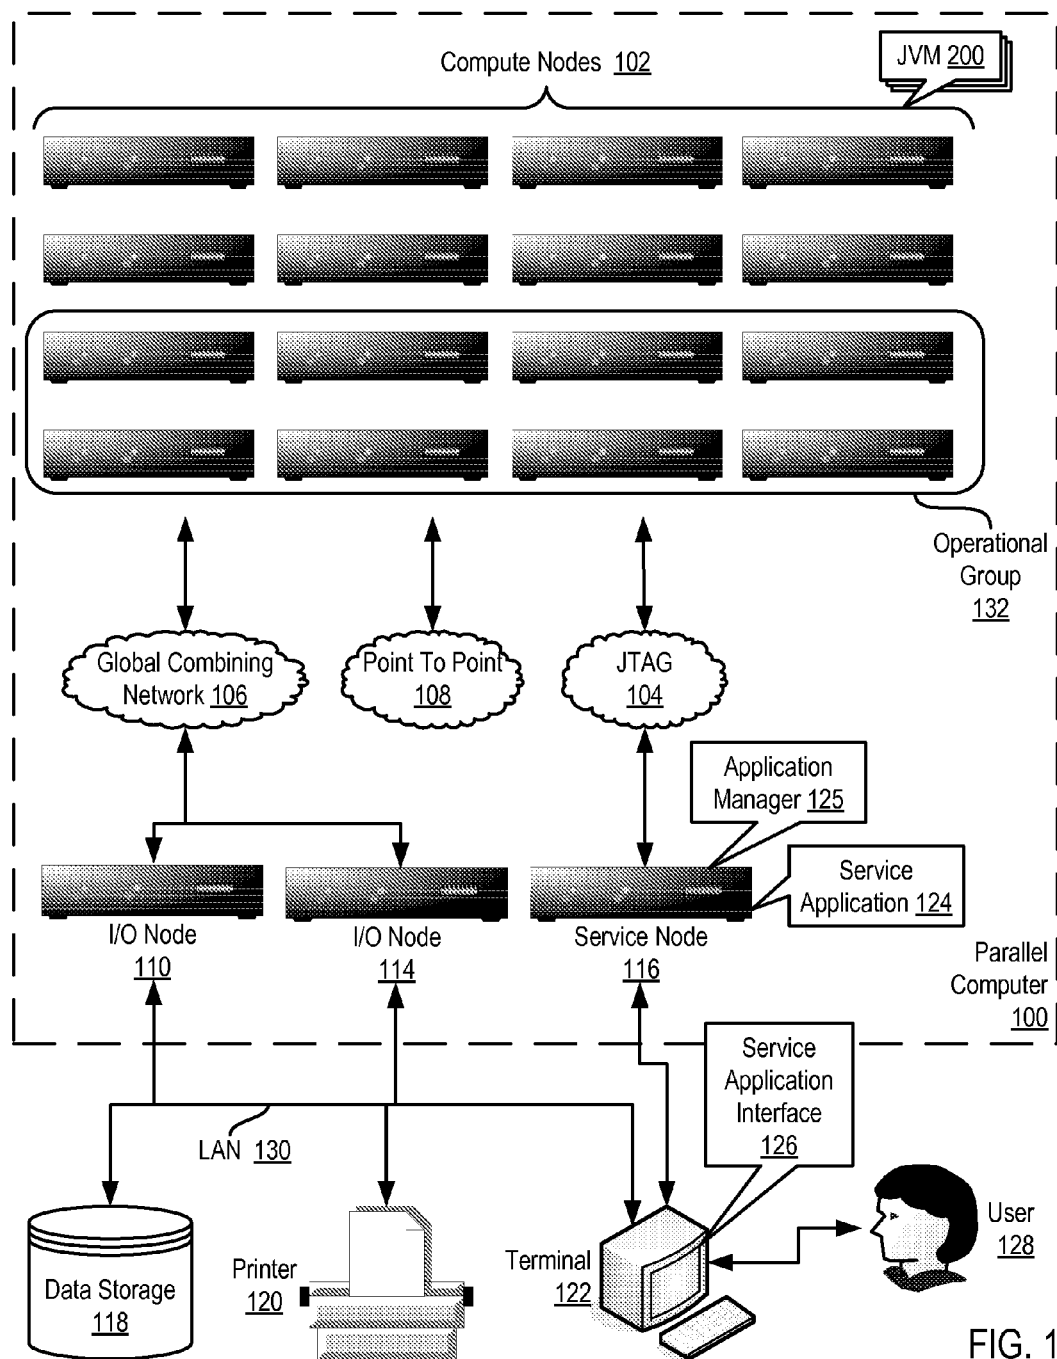
(74) Attorney, Agent, or Firm — Kennedy Lenart Spraggins LLP

#### (57) ABSTRACT

Methods, systems, and products are disclosed for executing a distributed JAVA application on a plurality of compute nodes. The JAVA application includes a plurality of jobs distributed among the plurality of compute nodes. The plurality of compute nodes are connected together for data communications through a data communication network. Each of the plurality of compute nodes has installed upon it a JAVA Virtual Machine ('JVM') capable of supporting at least one job of the JAVA application. Executing a distributed JAVA application on a plurality of compute nodes includes: tracking, by an application manager, JVM environment variables for the JVMs installed on the plurality of compute nodes; and configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the JVM environment variables for the JVMs installed on the plurality of compute nodes.

**18 Claims, 8 Drawing Sheets**





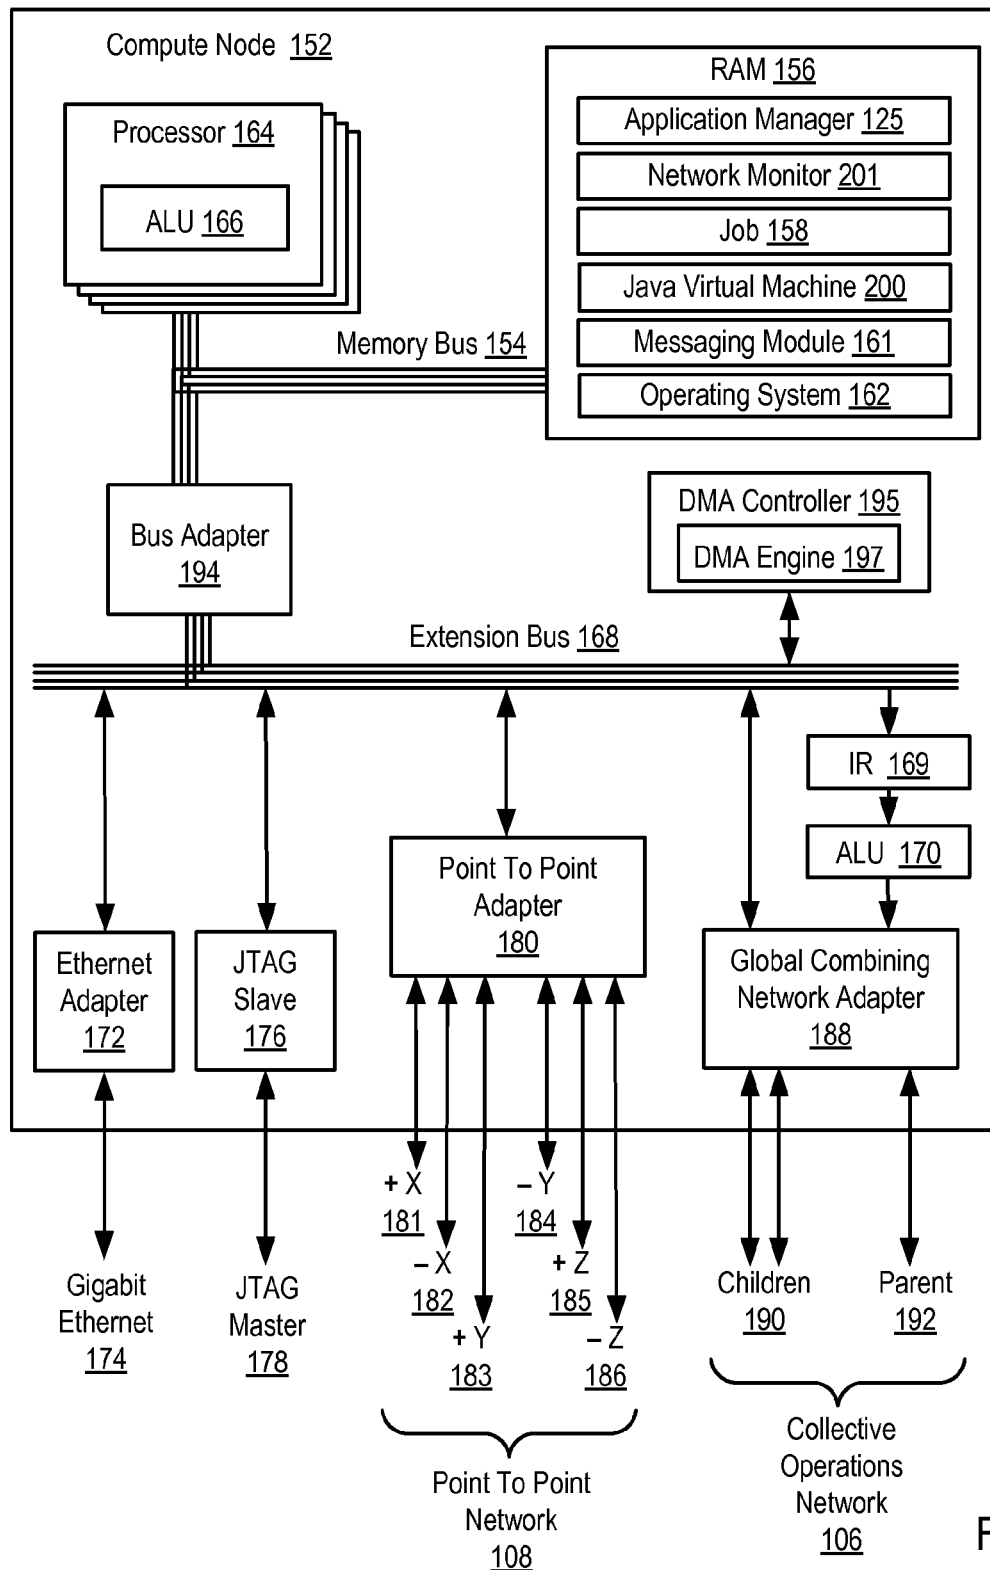


FIG. 2




FIG. 3A





FIG. 3B



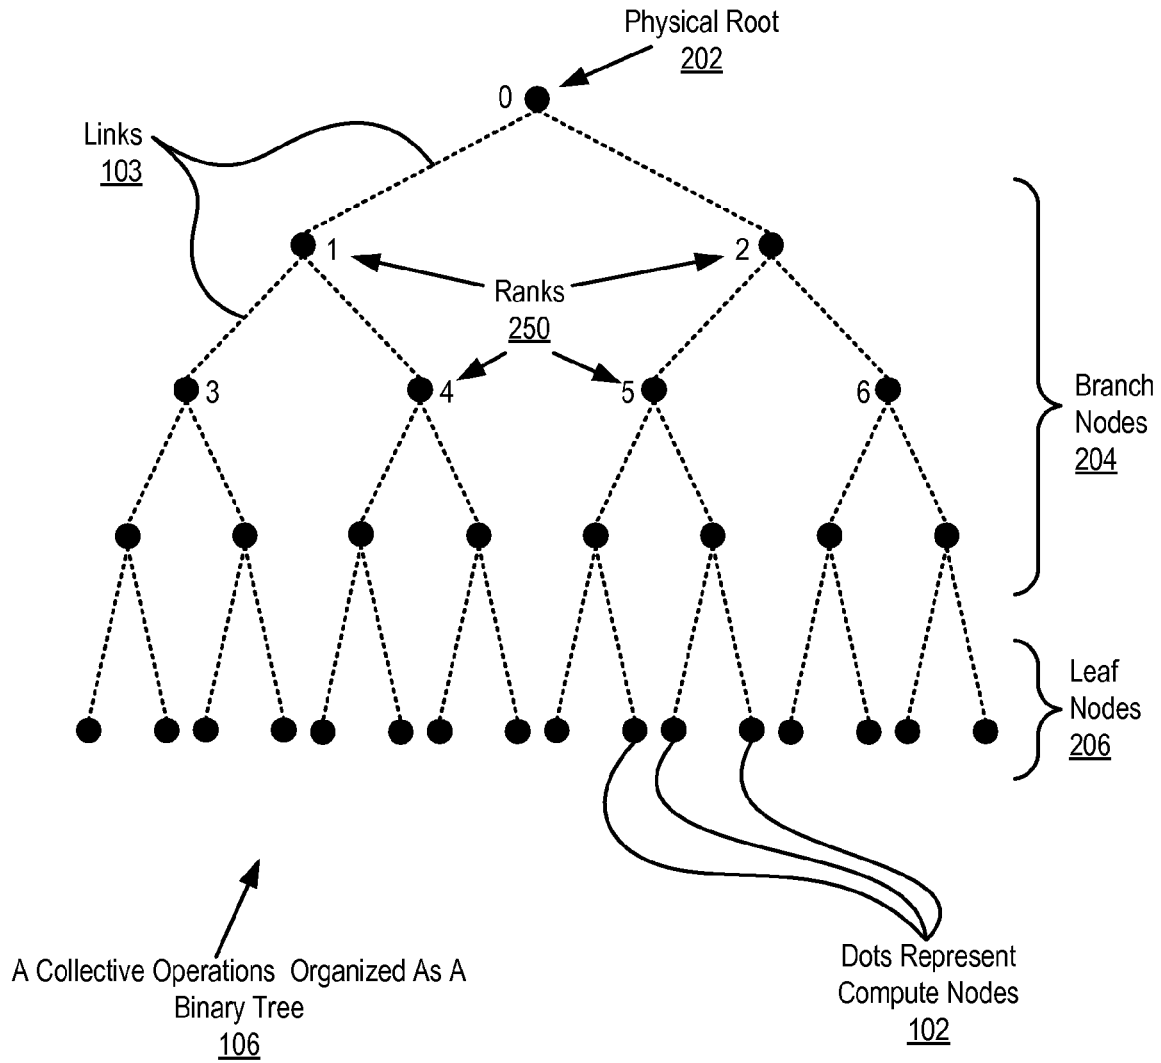


FIG. 5




FIG. 6

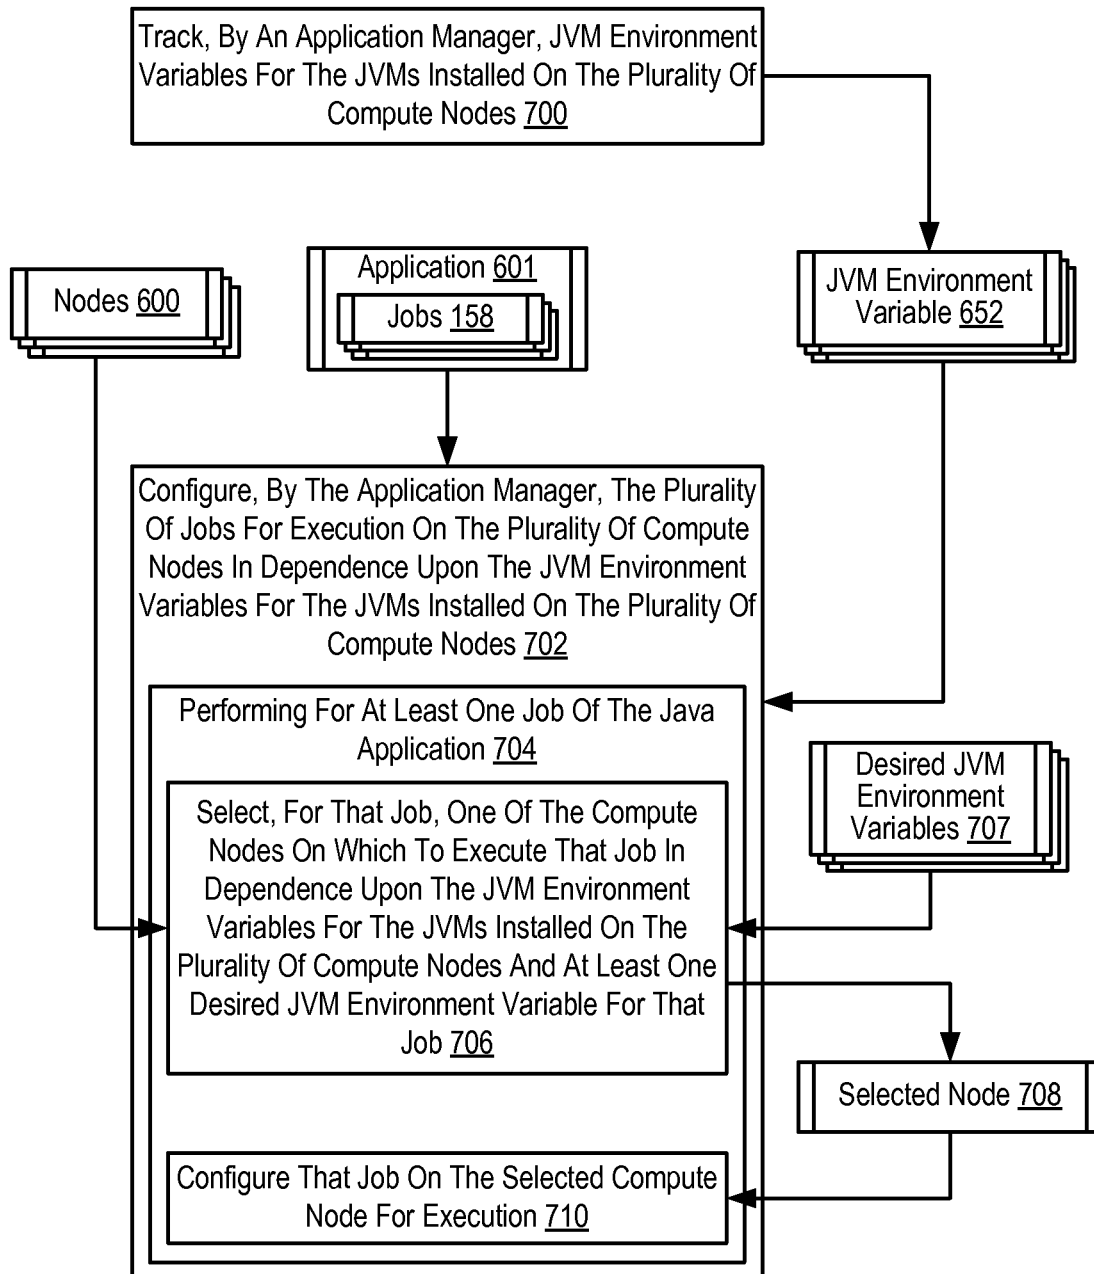


FIG. 7



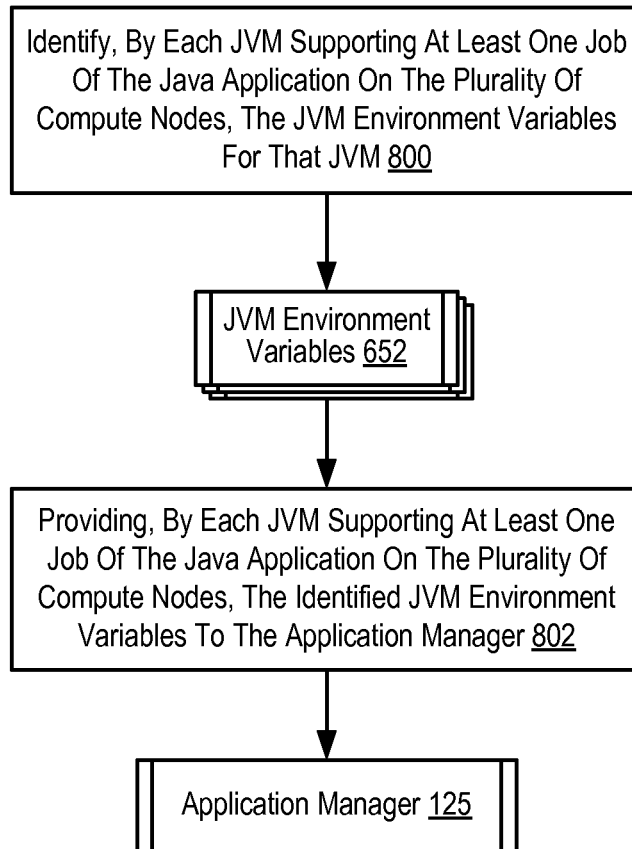


FIG. 8

1

## EXECUTING A DISTRIBUTED JAVA APPLICATION ON A PLURALITY OF COMPUTE NODES

### BACKGROUND OF THE INVENTION

#### 1. Field of the Invention

The field of the invention is data processing, or, more specifically, methods, apparatus, and products for executing a distributed JAVA application on a plurality of compute nodes.

#### 2. Description of Related Art

The development of the EDVAC computer system of 1948 is often cited as the beginning of the computer era. Since that time, computer systems have evolved into extremely complicated devices. Today's computers are much more sophisticated than early systems such as the EDVAC. Computer systems typically include a combination of hardware and software components, application programs, operating systems, processors, buses, memory, input/output devices, and so on. As advances in semiconductor processing and computer architecture push the performance of the computer higher and higher, more sophisticated computer software has evolved to take advantage of the higher performance of the hardware, resulting in computer systems today that are much more powerful than just a few years ago.

Parallel computing is an area of computer technology that has experienced advances. Parallel computing is the simultaneous execution of the same task (split up and specially adapted) on multiple processors in order to obtain results faster. Parallel computing is based on the fact that the process of solving a problem usually can be divided into smaller tasks, which may be carried out simultaneously with some coordination.

Parallel computers execute parallel algorithms. A parallel algorithm can be split up to be executed a piece at a time on many different processing devices, and then put back together again at the end to get a data processing result. Some algorithms are easy to divide up into pieces. Splitting up the job of checking all of the numbers from one to a hundred thousand to see which are primes could be done, for example, by assigning a subset of the numbers to each available processor, and then putting the list of positive results back together. In this specification, the multiple processing devices that execute the individual pieces of a parallel program are referred to as 'compute nodes.' A parallel computer is composed of compute nodes and other processing nodes as well, including, for example, input/output ('I/O') nodes, and service nodes.

Parallel algorithms are valuable because it is faster to perform some kinds of large computing tasks via a parallel algorithm than it is via a serial (non-parallel) algorithm, because of the way modern processors work. It is far more difficult to construct a computer with a single fast processor than one with many slow processors with the same throughput. There are also certain theoretical limits to the potential speed of serial processors. On the other hand, every parallel algorithm has a serial part and so parallel algorithms have a saturation point. After that point adding more processors does not yield any more throughput but only increases the overhead and cost.

Parallel algorithms are designed also to optimize one more resource the data communications requirements among the nodes of a parallel computer. There are two ways parallel processors communicate, shared memory or message passing. Shared memory processing needs additional locking for the data and imposes the overhead of additional processor and bus cycles and also serializes some portion of the algorithm.

2

Message passing processing uses high-speed data communications networks and message buffers, but this communication adds transfer overhead on the data communications networks as well as additional memory need for message buffers and latency in the data communications among nodes. Designs of parallel computers use specially designed data communications links so that the communication overhead will be small but it is the parallel algorithm that decides the volume of the traffic.

Many data communications network architectures are used for message passing among nodes in parallel computers. Compute nodes may be organized in a network as a 'torus' or 'mesh,' for example. Also, compute nodes may be organized in a network as a tree. A torus network connects the nodes in a three-dimensional mesh with wrap around links. Every node is connected to its six neighbors through this torus network, and each node is addressed by its x,y,z coordinate in the mesh. A torus network lends itself to point to point operations. In a tree network, the nodes typically are connected into a binary tree: each node has a parent, and two children (although some nodes may only have zero children or one child, depending on the hardware configuration). In computers that use a torus and a tree network, the two networks typically are implemented independently of one another, with separate routing circuits, separate physical links, and separate message buffers. A tree network provides high bandwidth and low latency for certain collective operations, message passing operations where all compute nodes participate simultaneously, such as, for example, an allgather.

The parallel applications that execute on the nodes in the data communications networks may be implemented in a variety of software programming languages, including the various versions and derivatives of JAVA™ technology promulgated by Sun Microsystems. JAVA applications generally run in a virtual execution environment called the JAVA Virtual Machine ('JVM'), rather than running directly on the computer hardware. The JAVA application is typically compiled into byte-code form, and then compiled in a just-in-time ('JIT') manner, or on-the-fly, by the JVM into JIT code representing hardware commands specific to the hardware platform on which the JVM is installed.

In a parallel computer, the JAVA application is generally a distributed application that is composed of multiple jobs; each job is typically implemented using one or more JAVA classes. Because the jobs are typically designed in a modular fashion, each job may be utilized in more than one JAVA application. The JVMs on the compute nodes of the parallel computer provide an execution environment for the jobs that make up a JAVA application. The execution performance of a particular job may vary from one JVM to another because each JVM may have different execution environment settings. Accordingly, executing a job on a compute node with a JVM that is not optimized for a particular job wastes valuable computing resources.

### SUMMARY OF THE INVENTION

Methods, systems, and products are disclosed for executing a distributed JAVA application on a plurality of compute nodes. The JAVA application includes a plurality of jobs distributed among the plurality of compute nodes. The plurality of compute nodes are connected together for data communications through a data communication network. Each of the plurality of compute nodes has installed upon it a JAVA Virtual Machine ('JVM') capable of supporting at least one job of the JAVA application. Executing a distributed JAVA application on a plurality of compute nodes includes: track-

ing, by an application manager, JVM environment variables for the JVMs installed on the plurality of compute nodes; and configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the JVM environment variables for the JVMs installed on the plurality of compute nodes.

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular descriptions of exemplary embodiments of the invention as illustrated in the accompanying drawings wherein like reference numbers generally represent like parts of exemplary embodiments of the invention.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an exemplary system for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

FIG. 2 sets forth a block diagram of an exemplary compute node useful in a parallel computer capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

FIG. 3A illustrates an exemplary Point To Point Adapter useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

FIG. 3B illustrates an exemplary Global Combining Network Adapter useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

FIG. 4 sets forth a line drawing illustrating an exemplary data communications network optimized for point to point operations useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes in accordance with embodiments of the present invention.

FIG. 5 sets forth a line drawing illustrating an exemplary data communications network optimized for collective operations useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes in accordance with embodiments of the present invention.

FIG. 6 sets forth a block diagram illustrating an exemplary system useful in executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

FIG. 7 sets forth a flow chart illustrating an exemplary method for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

FIG. 8 sets forth a flow chart illustrating a further exemplary method for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

### DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

Exemplary methods, apparatus, and computer program products for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention are described with reference to the accompanying drawings, beginning with FIG. 1. FIG. 1 illustrates an exemplary system for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. The system of FIG. 1 includes a parallel computer (100), non-volatile memory for the computer in the form of data storage device (118), an output device for the computer in the form of printer (120), and an

input/output device for the computer in the form of computer terminal (122). Parallel computer (100) in the example of FIG. 1 includes a plurality of compute nodes (102).

The compute nodes (102) are coupled for data communications by several independent data communications networks including a Joint Test Action Group ('JTAG') network (104), a global combining network (106) which is optimized for collective operations, and a torus network (108) which is optimized point to point operations. The global combining network (106) is a data communications network that includes data communications links connected to the compute nodes so as to organize the compute nodes as a tree. Each data communications network is implemented with data communications links among the compute nodes (102). The data communications links provide data communications for parallel operations among the compute nodes of the parallel computer. The links between compute nodes are bi-directional links that are typically implemented using two separate directional data communications paths.

In addition, the compute nodes (102) of parallel computer are organized into at least one operational group (132) of compute nodes for collective parallel operations on parallel computer (100). An operational group of compute nodes is the set of compute nodes upon which a collective parallel operation executes. Collective operations are implemented with data communications among the compute nodes of an operational group. Collective operations are those functions that involve all the compute nodes of an operational group. A collective operation is an operation, a message-passing computer program instruction that is executed simultaneously, that is, at approximately the same time, by all the compute nodes in an operational group of compute nodes. Such an operational group may include all the compute nodes in a parallel computer (100) or a subset all the compute nodes. Collective operations are often built around point to point operations. A collective operation requires that all processes on all compute nodes within an operational group call the same collective operation with matching arguments. A 'broadcast' is an example of a collective operation for moving data among compute nodes of an operational group. A 'reduce' operation is an example of a collective operation that executes arithmetic or logical functions on data distributed among the compute nodes of an operational group. An operational group may be implemented as, for example, an MPI 'communicator.'

'MPI' refers to 'Message Passing Interface,' a prior art parallel communications library, a module of computer program instructions for data communications on parallel computers. Examples of prior-art parallel communications libraries that may be improved for use with systems according to embodiments of the present invention include MPI and the 'Parallel Virtual Machine' ('PVM') library. PVM was developed by the University of Tennessee, The Oak Ridge National Laboratory, and Emory University. MPI is promulgated by the MPI Forum, an open group with representatives from many organizations that define and maintain the MPI standard. MPI at the time of this writing is a de facto standard for communication among compute nodes running a parallel program on a distributed memory parallel computer. This specification sometimes uses MPI terminology for ease of explanation, although the use of MPI as such is not a requirement or limitation of the present invention.

Some collective operations have a single originating or receiving process running on a particular compute node in an operational group. For example, in a 'broadcast' collective operation, the process on the compute node that distributes the data to all the other compute nodes is an originating

process. In a 'gather' operation, for example, the process on the compute node that received all the data from the other compute nodes is a receiving process. The compute node on which such an originating or receiving process runs is referred to as a logical root.

Most collective operations are variations or combinations of four basic operations: broadcast, gather, scatter, and reduce. The interfaces for these collective operations are defined in the MPI standards promulgated by the MPI Forum. Algorithms for executing collective operations, however, are not defined in the MPI standards. In a broadcast operation, all processes specify the same root process, whose buffer contents will be sent. Processes other than the root specify receive buffers. After the operation, all buffers contain the message from the root process.

In a scatter operation, the logical root divides data on the root into segments and distributes a different segment to each compute node in the operational group. In scatter operation, all processes typically specify the same receive count. The send arguments are only significant to the root process, whose buffer actually contains sendcount\*N elements of a given data type, where N is the number of processes in the given group of compute nodes. The send buffer is divided and dispersed to all processes (including the process on the logical root). Each compute node is assigned a sequential identifier termed a 'rank.' After the operation, the root has sent sendcount data elements to each process in increasing rank order. Rank 0 receives the first sendcount data elements from the send buffer. Rank 1 receives the second sendcount data elements from the send buffer, and so on.

A gather operation is a many-to-one collective operation that is a complete reverse of the description of the scatter operation. That is, a gather is a many-to-one collective operation in which elements of a datatype are gathered from the ranked compute nodes into a receive buffer in a root node.

A reduce operation is also a many-to-one collective operation that includes an arithmetic or logical function performed on two data elements. All processes specify the same 'count' and the same arithmetic or logical function. After the reduction, all processes have sent count data elements from computer node send buffers to the root process. In a reduction operation, data elements from corresponding send buffer locations are combined pair-wise by arithmetic or logical operations to yield a single corresponding element in the root process's receive buffer. Application specific reduction operations can be defined at runtime. Parallel communications libraries may support predefined operations. MPI, for example, provides the following pre-defined reduction operations:

MPI\_MAX maximum  
MPI\_MIN minimum  
MPI\_SUM sum  
MPI\_PROD product  
MPI\_LAND logical and  
MPI\_BAND bitwise and  
MPI\_LOR logical or  
MPI\_BOR bitwise or  
MPI\_LXOR logical exclusive or  
MPI\_BXOR bitwise exclusive or

In addition to compute nodes, the parallel computer (100) includes input/output ('I/O') nodes (110, 114) coupled to compute nodes (102) through the global combining network (106). The compute nodes in the parallel computer (100) are partitioned into processing sets such that each compute node in a processing set is connected for data communications to the same I/O node. Each processing set, therefore, is composed of one I/O node and a subset of compute nodes (102).

The ratio between the number of compute nodes to the number of I/O nodes in the entire system typically depends on the hardware configuration for the parallel computer. For example, in some configurations, each processing set may be composed of eight compute nodes and one I/O node. In some other configurations, each processing set may be composed of sixty-four compute nodes and one I/O node. Such example are for explanation only, however, and not for limitation. Each I/O nodes provide I/O services between compute nodes (102) of its processing set and a set of I/O devices. In the example of FIG. 1, the I/O nodes (110, 114) are connected for data communications I/O devices (118, 120, 122) through local area network (LAN) (130) implemented using high-speed Ethernet.

The parallel computer (100) of FIG. 1 also includes a service node (116) coupled to the compute nodes through one of the networks (104). Service node (116) provides services common to pluralities of compute nodes, administering the configuration of compute nodes, loading programs into the compute nodes, starting program execution on the compute nodes, retrieving results of program operations on the computer nodes, and so on. Service node (116) runs a service application (124) and communicates with users (128) through a service application interface (126) that runs on computer terminal (122).

In the example of FIG. 1, the service node (116) has installed upon it an application manager (125). The application manager (125) of FIG. 1 includes a set of computer program instructions capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. The JAVA application includes a plurality of jobs distributed among the plurality of compute nodes (102) for execution. The application manager (125) operates generally for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention by: tracking JVM environment variables for the JVMs installed on the plurality of compute nodes (102); and configuring the plurality of jobs for execution on the plurality of compute nodes (102) in dependence upon the JVM environment variables for the JVMs installed on the plurality of compute nodes (102). Although FIG. 1 illustrates the application manager (125) installed on a service node, readers will note that such an example is for explanation only and not for limitation. An application manager is a software component that may be installed on any compute nodes or other computer as will occur to those of skill in the art.

JVM environment variables represent settings for a particular JVM that affect the JVM's operation. JVM environment variables may include a JVM's minimum and maximum heap sizes, page size options, debug trace status, and so on. A JVM's heap is a storage area for the JAVA objects instantiated from JAVA classes, which form a typical job. Page size refers to the size of a contiguous block of virtual memory that is generally mapped to a frame in a page table. A frame is a contiguous block of physical memory used to store information. Debug trace is a debugging tool for a JVM that logs the execution of a set of computer program instructions.

Each compute node (102) of FIG. 1 has installed upon it a JAVA Virtual Machine ('JVM') (200) capable of supporting a JAVA application. Each JVM (200) of FIG. 1 includes a set of computer program instructions capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. Each JVM (200) operates generally for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention by: identifying the

JVM environment variables for that JVM (200); and providing the identified JVM environment variables to the application manager (125).

In the example of FIG. 1, the plurality of compute nodes (102) are implemented in a parallel computer (100) and are connected together using a plurality of data communications networks (104, 106, 108). The point to point network (108) is optimized for point to point operations. The global combining network (106) is optimized for collective operations. Although executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention is described above in terms of executing a distributed JAVA application on a parallel computer, readers will note that such an embodiment is for explanation only and not for limitation. In fact, executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention may be implemented using a variety of computer systems composed of a plurality of nodes network-connected together, including for example a cluster of nodes, a distributed computing system, a grid computing system, and so on.

The arrangement of nodes, networks, and I/O devices making up the exemplary system illustrated in FIG. 1 are for explanation only, not for limitation of the present invention. Data processing systems capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention may include additional nodes, networks, devices, and architectures, not shown in FIG. 1, as will occur to those of skill in the art. Although the parallel computer (100) in the example of FIG. 1 includes sixteen compute nodes (102), readers will note that parallel computers capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention may include any number of compute nodes. In addition to Ethernet and JTAG, networks in such data processing systems may support many data communications protocols including for example TCP (Transmission Control Protocol), IP (Internet Protocol), and others as will occur to those of skill in the art. Various embodiments of the present invention may be implemented on a variety of hardware platforms in addition to those illustrated in FIG. 1.

Executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention may be generally implemented on a parallel computer that includes a plurality of compute nodes, among other types of exemplary systems. In fact, such computers may include thousands of such compute nodes. Each compute node is in turn itself a kind of computer composed of one or more computer processors, its own computer memory, and its own input/output adapters. For further explanation, therefore, FIG. 2 sets forth a block diagram of an exemplary compute node (152) useful in a parallel computer capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

The compute node (152) of FIG. 2 includes one or more computer processors (164) as well as random access memory ('RAM') (156). The processors (164) are connected to RAM (156) through a high-speed memory bus (154) and through a bus adapter (194) and an extension bus (168) to other components of the compute node (152). Stored in RAM (156) is a job (158) of a JAVA application distributed across the nodes of the parallel computer. The job (158) is a module of computer program instructions, typically represented in byte code, that carries out parallel, user-level data processing using one or more JAVA classes.

Also stored in RAM (156) is an application manager (125). The application manager (125) of FIG. 2 includes a set of

computer program instructions capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. The application manager (125) operates generally for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention by: tracking JVM environment variables for the JVMs installed on the plurality of compute nodes and configuring the plurality of jobs for execution on the plurality of compute nodes in dependence upon the JVM environment variables for the JVMs installed on the plurality of compute nodes.

Also stored in RAM (156) is a network monitor (201) that monitors the utilization of each of the nodes and the data communication networks connecting the nodes together. The network monitors on the compute nodes communicate with one another to keep all of the nodes and any service nodes informed about the state of the individual nodes and the networks. In addition, the network monitor (201) may be used to facilitate the exchange of JVM environment variables between the JVMs installed on the compute nodes and the application manager (125). To facilitate such communications each network monitor may register with the application manager (125) and expose and interface to the JVM through which the JVM can update the application manager (125) with the JVM's environment variables. In such a manner, the network monitor (201) may serve a 'local client' for the application manager (125) on each of the nodes.

Also stored in RAM (156) is a JVM (200). The JVM (200) of FIG. 2 is a set of computer software programs and data structures which implements a virtual execution environment for a specific hardware platform. The JVM (200) of FIG. 2 accepts the JAVA application (158) for execution in a computer intermediate language, commonly referred to as JAVA byte code, which is a hardware-independent compiled form of the JAVA application (158). In such a manner, the JVM (200) of FIG. 2 serves to abstract the compiled version of the JAVA application (158) from the hardware of node (152) because the JVM (200) handles the hardware specific implementation details of executing the application (158) during runtime. Abstracting the hardware details of a platform from the compiled form of a JAVA application allows the application to be compiled once into byte code, yet run on a variety of hardware platforms.

The JVM (200) of FIG. 2 is improved for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. The JVM (200) of FIG. 2 operates generally for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention by: identifying the JVM environment variables for that JVM (200) and providing the identified JVM environment variables to the application manager (125).

Also stored RAM (156) is a messaging module (161), a library of computer program instructions that carry out parallel communications among compute nodes, including point to point operations as well as collective operations. The JAVA application (158) effects data communications with other applications running on other compute nodes by calling software routines in the messaging modules (161). A library of parallel communications routines may be developed from scratch for use in systems according to embodiments of the present invention, using a traditional programming language such as the C programming language, and using traditional programming methods to write parallel communications routines. Alternatively, existing prior art libraries may be used such as, for example, the 'Message Passing Interface' ('MPI')

library, the 'Parallel Virtual Machine' ('PVM') library, and the Aggregate Remote Memory Copy Interface ('ARMC') library.

Also stored in RAM (156) is an operating system (162), a module of computer program instructions and routines for an application program's access to other resources of the compute node. It is typical for an application program and parallel communications library in a compute node of a parallel computer to run a single thread of execution with no user login and no security issues because the thread is entitled to complete access to all resources of the node. The quantity and complexity of tasks to be performed by an operating system on a compute node in a parallel computer therefore are smaller and less complex than those of an operating system on a serial computer with many threads running simultaneously. In addition, there is no video I/O on the compute node (152) of FIG. 2, another factor that decreases the demands on the operating system. The operating system may therefore be quite lightweight by comparison with operating systems of general purpose computers, a pared down version as it were, or an operating system developed specifically for operations on a particular parallel computer. Operating systems that may usefully be improved, simplified, for use in a compute node include UNIX™, Linux™, Microsoft Vista™, AIX™, IBM's i5/OS™, and others as will occur to those of skill in the art.

The exemplary compute node (152) of FIG. 2 includes several communications adapters (172, 176, 180, 188) for implementing data communications with other nodes of a parallel computer. Such data communications may be carried out serially through RS-232 connections, through external buses such as USB, through data communications networks such as IP networks, and in other ways as will occur to those of skill in the art. Communications adapters implement the hardware level of data communications through which one computer sends data communications to another computer, directly or through a network. Examples of communications adapters useful in systems for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention include modems for wired communications, Ethernet (IEEE 802.3) adapters for wired network communications, and 802.11b adapters for wireless network communications.

The data communications adapters in the example of FIG. 2 include a Gigabit Ethernet adapter (172) that couples example compute node (152) for data communications to a Gigabit Ethernet (174). Gigabit Ethernet is a network transmission standard, defined in the IEEE 802.3 standard, that provides a data rate of 1 billion bits per second (one gigabit). Gigabit Ethernet is a variant of Ethernet that operates over multimode fiber optic cable, single mode fiber optic cable, or unshielded twisted pair.

The data communications adapters in the example of FIG. 2 includes a JTAG Slave circuit (176) that couples example compute node (152) for data communications to a JTAG Master circuit (178). JTAG is the usual name used for the IEEE 1149.1 standard entitled Standard Test Access Port and Boundary-Scan Architecture for test access ports used for testing printed circuit boards using boundary scan. JTAG is so widely adapted that, at this time, boundary scan is more or less synonymous with JTAG. JTAG is used not only for printed circuit boards, but also for conducting boundary scans of integrated circuits, and is also useful as a mechanism for debugging embedded systems, providing a convenient "back door" into the system. The example compute node of FIG. 2 may be all three of these: It typically includes one or more integrated circuits installed on a printed circuit board and may be implemented as an embedded system having its own pro-

cessor, its own memory, and its own I/O capability. JTAG boundary scans through JTAG Slave (176) may efficiently configure processor registers and memory in compute node (152) for use in executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

The data communications adapters in the example of FIG. 2 includes a Point To Point Adapter (180) that couples example compute node (152) for data communications to a network (108) that is optimal for point to point message passing operations such as, for example, a network configured as a three-dimensional torus or mesh. Point To Point Adapter (180) provides data communications in six directions on three communications axes, x, y, and z, through six bidirectional links: +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186).

The data communications adapters in the example of FIG. 2 includes a Global Combining Network Adapter (188) that couples example compute node (152) for data communications to a network (106) that is optimal for collective message passing operations on a global combining network configured, for example, as a binary tree. The Global Combining Network Adapter (188) provides data communications through three bidirectional links: two to children nodes (190) and one to a parent node (192).

Example compute node (152) includes two arithmetic logic units ('ALUs'). ALU (166) is a component of processor (164), and a separate ALU (170) is dedicated to the exclusive use of Global Combining Network Adapter (188) for use in performing the arithmetic and logical functions of reduction operations. Computer program instructions of a reduction routine in parallel communications library (160) may latch an instruction for an arithmetic or logical function into instruction register (169). When the arithmetic or logical function of a reduction operation is a 'sum' or a 'logical or,' for example, Global Combining Network Adapter (188) may execute the arithmetic or logical operation by use of ALU (166) in processor (164) or, typically much faster, by use dedicated ALU (170).

The example compute node (152) of FIG. 2 includes a direct memory access ('DMA') controller (195), which is computer hardware for direct memory access and a DMA engine (195), which is computer software for direct memory access. Direct memory access includes reading and writing to memory of compute nodes with reduced operational burden on the central processing units (164). A DMA transfer essentially copies a block of memory from one compute node to another. While the CPU may initiate the DMA transfer, the CPU does not execute it. In the example of FIG. 2, the DMA engine (195) and the DMA controller (195) support the messaging module (161).

For further explanation, FIG. 3A illustrates an exemplary Point To Point Adapter (180) useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. Point To Point Adapter (180) is designed for use in a data communications network optimized for point to point operations, a network that organizes compute nodes in a three-dimensional torus or mesh. Point To Point Adapter (180) in the example of FIG. 3A provides data communication along an x-axis through four unidirectional data communications links, to and from the next node in the -x direction (182) and to and from the next node in the +x direction (181). Point To Point Adapter (180) also provides data communication along a y-axis through four unidirectional data communications links, to and from the next node in the -y direction (184) and to and from the next node in the +y direction (183).

Point To Point Adapter (180) in FIG. 3A also provides data communication along a z-axis through four unidirectional data communications links, to and from the next node in the -z direction (186) and to and from the next node in the +z direction (185).

For further explanation, FIG. 3B illustrates an exemplary Global Combining Network Adapter (188) useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. Global Combining Network Adapter (188) is designed for use in a network optimized for collective operations, a network that organizes compute nodes of a parallel computer in a binary tree. Global Combining Network Adapter (188) in the example of FIG. 3B provides data communication to and from two children nodes through four unidirectional data communications links (190). Global Combining Network Adapter (188) also provides data communication to and from a parent node through two unidirectional data communications links (192).

For further explanation, FIG. 4 sets forth a line drawing illustrating an exemplary data communications network (108) optimized for point to point operations useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes in accordance with embodiments of the present invention. In the example of FIG. 4, dots represent compute nodes (102) of a parallel computer, and the dotted lines between the dots represent data communications links (103) between compute nodes. The data communications links are implemented with point to point data communications adapters similar to the one illustrated for example in FIG. 3A, with data communications links on three axes, x, y, and z, and to and from in six directions +x (181), -x (182), +y (183), -y (184), +z (185), and -z (186). The links and compute nodes are organized by this data communications network optimized for point to point operations into a three dimensional mesh (105). The mesh (105) has wrap-around links on each axis that connect the outermost compute nodes in the mesh (105) on opposite sides of the mesh (105). These wrap-around links form part of a torus (107). Each compute node in the torus has a location in the torus that is uniquely specified by a set of x, y, z coordinates. Readers will note that the wrap-around links in the y and z directions have been omitted for clarity, but are configured in a similar manner to the wrap-around link illustrated in the x direction. For clarity of explanation, the data communications network of FIG. 4 is illustrated with only 27 compute nodes, but readers will recognize that a data communications network optimized for point to point operations for use in executing a distributed JAVA application on a plurality of compute nodes in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes.

For further explanation, FIG. 5 sets forth a line drawing illustrating an exemplary data communications network (106) optimized for collective operations useful in systems capable of executing a distributed JAVA application on a plurality of compute nodes in accordance with embodiments of the present invention. The example data communications network of FIG. 5 includes data communications links connected to the compute nodes so as to organize the compute nodes as a tree. In the example of FIG. 5, dots represent compute nodes (102) of a parallel computer, and the dotted lines (103) between the dots represent data communications links between compute nodes. The data communications links are implemented with global combining network adapters similar to the one illustrated for example in FIG. 3B, with each node typically providing data communications to and

from two children nodes and data communications to and from a parent node, with some exceptions. Nodes in a binary tree (106) may be characterized as a physical root node (202), branch nodes (204), and leaf nodes (206). The root node (202) has two children but no parent. The leaf nodes (206) each has a parent, but leaf nodes have no children. The branch nodes (204) each has both a parent and two children. The links and compute nodes are thereby organized by this data communications network optimized for collective operations into a binary tree (106). For clarity of explanation, the data communications network of FIG. 5 is illustrated with only 31 compute nodes, but readers will recognize that a data communications network optimized for collective operations for use in systems for executing a distributed JAVA application on a plurality of compute nodes in accordance with embodiments of the present invention may contain only a few compute nodes or may contain thousands of compute nodes.

In the example of FIG. 5, each node in the tree is assigned a unit identifier referred to as a 'rank' (250). A node's rank uniquely identifies the node's location in the tree network for use in both point to point and collective operations in the tree network. The ranks in this example are assigned as integers beginning with 0 assigned to the root node (202), 1 assigned to the first node in the second layer of the tree, 2 assigned to the second node in the second layer of the tree, 3 assigned to the first node in the third layer of the tree, 4 assigned to the second node in the third layer of the tree, and so on. For ease of illustration, only the ranks of the first three layers of the tree are shown here, but all compute nodes in the tree network are assigned a unique rank.

For further explanation, FIG. 6 sets forth a block diagram illustrating an exemplary system useful in executing a distributed JAVA application (601) on a plurality of compute nodes (600a) according to embodiments of the present invention. The JAVA application (601) of FIG. 6 includes a plurality of jobs (158) distributed among the plurality of compute nodes (600). Each compute node (600) has installed upon it a JVM (200) capable of supporting at least one job (158) of the JAVA application (601).

The nodes (600) of FIG. 6 are connected together for data communications using a data communication network. In addition, the nodes (600) are connected to an I/O node (110) that provides I/O services between the nodes (600) and a set of I/O devices such as, for example, the service node (116) and the data storage (118). The service node (116) of FIG. 6 provides services common to nodes (600), administering the configuration of nodes (600), loading programs such as JAVA application (601) and JVM (200) onto the nodes (600), starting program execution on the nodes (600), retrieving results of program operations on the nodes (600), and so on. The data storage (118) of FIG. 6 may store the files that contain the JAVA classes that compose the JAVA application (601).

The service node (116) has installed upon it an application manager (125). The application manager (125) includes a set of computer program instructions capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. The application manager (125) operates generally for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention by: tracking JVM environment variables (652) for the JVMs (200) installed on the plurality of compute nodes (600) and configuring the plurality of jobs (158) for execution on the plurality of compute nodes (600) in dependence upon the JVM environment variables (652) for the JVMs (200) installed on the plurality of compute nodes (600). In such a manner, the application manager (125) may configure a job (158) of the appli-

## 13

cation (601) for execution on a compute node (600) that has JVM environment variables that match optimal or desired JVM environment variables specified by an application profile (603) for the application (601).

In the example of FIG. 6, the application profile (603) of FIG. 6 is a data structure used by the application manager (125) that associates desired JVM environment variables with various jobs (158) of the application (601). The application profile (603) may be formatted as a text file, a table, a structured document, or any other format as will occur to those of skill in the art. In addition to using the application profile (603) to configure jobs (158) on the nodes (600), the application manager (125) may also be used to assign desired JVM environment variables to the various jobs (158) of application (601) in the application profile (603). The application manager (125) may assign desired JVM environment variables to the various jobs (158) of application (601) based on user-specified environmental settings for the various jobs (158) of the application (601) or based on a historic execution performance for the various jobs (158) of the application (601).

In the example of FIG. 6, the application manager (125) configures job (158a) on compute node (600a) for execution based on the JVM environment variables (652a) for the JVM (200) on node (600a). To execute the job (158a), the JVM (200) of FIG. 6 includes a JIT compiler (618) and JIT code executor (654). The JIT compiler (618) compiles the byte code representation of job (158a) into machine code to be run directly on the native platform and optimizes the machine code for enhanced performance. The JIT compiler (618) is typically invoked when the job (158a) is started up or when some other usage criterion is met to improve run-time performance by avoiding the need for this code to be interpreted later. The output from the JIT compiler (618) is JIT code (616), which is in turn provided into the JIT code executor (654) for execution on the platform's processor.

The JIT compiler (618) operates by spawning a compilation thread that receives the byte-code version of the job (158a), translates and optimizes the byte code into native code, and feeds the native code to an execution thread for the JIT code executor (654). After processing various portions of the job (158a), the JIT compiler (618) examines the portions of the job (158a) executed by the JIT code executor (654). The JIT compiler (618) may monitor how the JIT code executor (654) executed the job (158a) to determine any further optimizations that would enhance future execution performance. For example, if portions of the job (158a) are consistently being skipped during execution, the JIT compiler (618) may not process those skipped portions in the future. Similarly, if portions of the job (158a) are being executed repeatedly, the JIT compiler (618) attempt to further optimize those heavily executed portions.

In the example of FIG. 6, the JVM (200) is improved for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention. The JVM (200) operates generally for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention by: identifying the JVM environment variables (652) for that JVM (200) and providing the identified JVM environment variables (652) to the application manager (125). The JVM (200) may communicate with the application manager (125) through the network monitor (201). The application manager (125) may then in turn utilize the JVM environment variables (652) received from each of the JVMs supporting the jobs (158) to configure the jobs (158) in the future on the nodes

## 14

(600) having a JVM with environment variables that match the desired, or optimal, JVM environment variables for those jobs (158).

The job (158a) of the application (601) often represents any number of JAVA classes. As the JVM (200) executes the job (158a), the JVM identifies a JAVA class utilized for the job (158a), loads the JAVA classes for the job (158a) into memory, and prepares each class instance for execution. The JVM (200) therefore includes a hierarchy of class loaders (620) that operate to load the classes specified by the job (158a). The hierarchy of class loaders (620) includes a primordial class loader (622), an extension class loader (624), and an application class loader (626).

The primordial class loader (622) of FIG. 6 loads the core JAVA libraries, such as 'core.jar,' 'server.jar,' and so on, in the '<JAVA\_HOME>/lib' directory. The primordial class loader (622), which is part of the core JVM, is written in native code specific to the hardware platform on which the JVM is installed. The extension class loader (624) of FIG. 6 loads the code in the extensions directories and is typically implemented by the 'sun.misc.Launcher\$ExtClassLoader' class. The application class loader (626) of FIG. 6 loads the class specified by 'java.class.path,' which maps to the system 'CLASSPATH' variable. The application class loader (626) is typically implemented by the 'sun.misc.Launcher\$AppClassLoader' class.

For each class included or specified by the job (158a), the JVM (200) effectively traverses up the class loader hierarchy to determine whether any class loader has previously loaded the class. The order of traversal is as follows: first to the default application class loader (626), then to the extension class loader (624), and finally to the primordial class loader (622). If the response from all of the class loaders is negative, then the JVM (200) traverses down the hierarchy, with the primordial class loader first attempting to locate the class by searching the locations specified in its class path definition. If the primordial class loader (622) is unsuccessful, then the then the extension class loader (624) may make a similar attempt to load the class. If the extension class loader (624) is unsuccessful, then the application class loader (626) attempts to load the class. Finally, if the application class loader (626) is unsuccessful, then the JVM (200) triggers an error condition.

The JVM (200) of FIG. 6 also includes a heap (610), which is shared between all threads, and is used for storage of objects (612). Each object (612) represents an already loaded class. That is, each object (612) is in effect an instantiation of a class, which defines the object. Because a job may utilize more than one object of the same type, a single class may be instantiated multiple times to create the objects specified by the job (158a). Readers will note that the class loaders (620) are objects that are also stored on heap (610), but for the sake of clarity the class loaders (620) are shown separately in FIG. 6.

In the example of FIG. 6, the JVM (200) also includes a class storage area (636), which is used for storing information relating to the classes stored in the heap (610). The class storage area (636) includes a method code region (638) for storing byte code for implementing class method calls, and a constant pool (640) for storing strings and other constants associated with a class. The class storage area (636) also includes a field data region (642) for sharing static variables, which are shared between all instances of a class, and a static initialization area (646) for storing static initialization methods and other specialized methods separate from the method code region (638). The class storage area also includes a method block area (644), which is used to stored information



15

relating to the code, such as invokers, and a pointer to the code, which may for example be in method code area (638), in JIT code area (616) described above, or loaded as native code such as, for example, a dynamic link library ('DLL') written in C or C++.

A class stored as an object (612) in the heap (610) contains a reference to its associated data, such as method byte code, in class storage area (636). Each object (612) contains a reference to the class loader (620), which loaded the class into the heap (610), plus other fields such as a flag to indicate whether or not they have been initialized.

In the example of FIG. 6, the JVM (200) also includes a stack area (614), which is used for storing the stacks associated with the execution of different threads on the JVM (200). Readers will note that because the system libraries and indeed parts of the JVM (200) itself are written in JAVA, which frequently utilize multi-threading, the JVM (200) may be supporting multiple threads even if the job (158a) contains only a single thread.

Also included within JVM (200) of FIG. 6 is a class loader cache (634) and garbage collector (650). The former is typically implemented as a table that allows a class loader to trace those classes which it initially loaded into the JVM (200). The class loader cache (634) therefore allows each class loader (620) to determine whether it has already loaded a particular class when the JVM (200) initially traverses the class loader hierarchy as described above. Readers will note that it is part of the overall security policy of the JVM (200) that classes will typically have different levels of permission within the system based on the identity of the class loader by which they were originally loaded.

The garbage collector (650) is used to delete objects (612) from heap (610) when they are no longer required. Thus in the JAVA programming language, applications do not need to specifically request or release memory, rather this is controlled by the JVM (200) itself. Therefore, when the job (158a) specifies the creation of an object (612), the JVM (200) secures the requisite memory resource. Then, when the job (158a) finishes using object (612), the JVM (200) can delete the object (612) to free up this memory resource. This process of deleting an object is known as 'garbage collection,' and is generally performed by briefly interrupting all threads on the stack (614), and scanning the heap (610) for objects (612) which are no longer referenced, and therefore can be deleted. The details of garbage collection vary from one JVM (200) implementation to another, but typically garbage collection is scheduled when the heap (610) is nearly exhausted and so there is a need to free up space for new objects (612).

In the example of FIG. 6, the JVM (200) also includes a monitor pool (648). The monitor pool (648) is used to store a set of locks or 'monitors' that are used to control contention to an object resulting from concurrent attempts to access the object by different threads when exclusive access to the object is required.

Although the JVM (200) in FIG. 6 is shown on and described above with regard to the node (600a), readers will note that each of the other nodes (600b) also has installed upon it a JVM configured in a similar manner. That is, each of the other nodes (600b) also has installed upon it a JVM capable of executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

For further explanation, FIG. 7 sets forth a flow chart illustrating an exemplary method for executing a distributed JAVA application (601) on a plurality of compute nodes (600) according to embodiments of the present invention. The JAVA application (601) includes a plurality of jobs (158) distributed

16

among the plurality of compute nodes (600). The plurality of compute nodes (600) connected together for data communications through a data communication network. Each of the plurality of compute nodes (600) has installed upon it a JVM capable of supporting at least one job (158) of the JAVA application (601).

The method of FIG. 7 includes tracking (700), by an application manager, a JVM environment variables (652) for the JVMs installed on the plurality of compute nodes. The application manager may track (700) a JVM environment variables (652) for the JVMs installed on the plurality of compute nodes according to the method of FIG. 7 by receiving the JVM environment variables (652) from each of the JVMs installed on the compute nodes (600). The application manager may receive the JVM environment variables (652) as each JVM transmit its JVM environment variables (652) periodically to the application manager. In some other embodiments, however, the application manager may periodically poll each JVM for the JVM environment variables (652). The application manager and the JVMs may communicate directly with one another or through some intermediary communications facilitator such as, for example, a network monitor.

The JVM environment variables (652) of FIG. 7 represent settings for a particular JVM that affect the JVM's operation. The JVM environment variables (652) may include a JVM's minimum and maximum heap sizes, page size options, debug trace status, and so on. For further explanation, consider the following exemplary JVM environment variables provided by one of the JVMs installed on the nodes (600):

TABLE 1

EXEMPLARY JVM ENVIRONMENT VARIABLES	
VARIABLE IDENTIFIER	VARIABLE VALUE
Minimum Heap Size	4 MB
Maximum Heap Size	1024 MB
Page Size	4 MB
Debug Trace Status	Disabled
Share Class Data	Enabled
JIT Optimization	Enabled
Verify Heap Integrity	Enabled
Background JIT Compilation	Enabled
...	...

The exemplary JVM environment variables in Table 1 above specify that the minimum heap size is four megabytes (MB) and the maximum heap size is one thousand twenty-four MB for the JVM. The exemplary JVM environment variables also specifies that JVM's page size is four MB and the Debug Trace Status for the JVM is disabled. The exemplary JVM environment variables specify that the JVM shares class data when possible and that JIT compilation optimizations are enabled. The exemplary JVM environment variables also specify that the JVM's heap integrity verification is enabled along with the JVM's option to perform JIT compilation in the background. Readers will note that the exemplary JVM environment variables above are for explanation only and not for limitation.

The method of FIG. 7 also includes configuring (702), by the application manager, the plurality of jobs (158) for execution on the plurality of compute nodes (600) in dependence upon the JVM environment variables (652) for the JVMs installed on the plurality of compute nodes (600). The application manager configures (702) the plurality of jobs (158) for execution on the plurality of compute nodes (600) according to the method of FIG. 7 by performing (704) for at least

17

one job (158) of the JAVA application (601): selecting (706), for that job, one of the compute nodes (600) on which to execute that job (158) in dependence upon the JVM environment variables (652) for the JVMs and at least one desired JVM environment variable (707) for that job (158); and configuring (710) that job (158) on the selected compute node (708) for execution.

The desired JVM environment variables (707) of FIG. 7 represent user-specified JVM environment variables for each job (158) of the JAVA application (601). The desired JVM environment variables (707) may be specified in an application profile for the application (601) that is provided by a system administrator or application developer. For example, consider the following exemplary application profile:

TABLE 2

EXEMPLARY APPLICATION PROFILE		
JOB IDENTIFIER	VARIABLE IDENTIFIER	VARIABLE VALUE
0	Minimum Heap Size	4 MB
0	Maximum Heap Size	1024 MB
0	Page Size	4 MB
0	Debug Trace Status	Disabled
0	Share Class Data	Enabled
0	JIT Optimization	Enabled
0	Verify Heap Integrity	Disabled
0	Background JIT Compilation	Enabled
1	...	...

The exemplary application profile in Table 2 above specifies that the desired minimum heap size is four MB and the desired maximum heap size is one thousand twenty-four MB for job 0. The exemplary application profile also specifies that desired page size for job 0 is four MB and the desired Debug Trace Status for the job 0 is disabled. The exemplary application profile specifies that, ideally for job 0, the JVM shares class data when possible and that JIT compilation optimizations are enabled. The exemplary application profile does specify that the JVM's heap integrity verification should be disabled for job 0, but that JVM's option to perform JIT compilation in the background should be enabled for job 0. Readers will note that the exemplary application profile above is for explanation only and not for limitation.

In the method of FIG. 7, the application manager may select (706) one of the compute nodes (600) on which to execute a particular job (158) by comparing the JVM environment variables (652) received from each of the JVMs with the desired JVM environment variables (707) for the particular job (158) and selecting the node (708) having the JVM environment variables (652) that best matches the desired JVM environment variables (707) for the particular job (158). The application manager may identify the best match as the compute node having the highest number of JVM environment variables (652) that match the desired JVM environment variables (707). For example, consider that the application manager may configure a job on node 0 or node 1. Further consider that node 0 provides the application manager with the exemplary JVM environment variables in Table 1 above and that node 1 provides the application manager with the following exemplary JVM environment variables:

18

TABLE 3

EXEMPLARY JVM ENVIRONMENT VARIABLES	
VARIABLE IDENTIFIER	VARIABLE VALUE
Minimum Heap Size	1 MB
Maximum Heap Size	512 MB
Page Size	1 MB
Debug Trace Status	Enabled
Share Class Data	Disabled
JIT Optimization	Disabled
Verify Heap Integrity	Disabled
Background JIT Compilation	Disabled
...	...

When comparing the desired JVM environment variables for job 0 in Table 2 with the exemplary JVM environment variables in Tables 1 for node 0 and in Table 3 for node 1, the application manager determines that node 0 has the highest number of JVM environment variables that match the desired JVM environment variables for job 0. Specifically, node 0 has seven JVM environment variables that match the desired JVM environment variables for job 0, while node 1 does not have any JVM environment variables that match the desired JVM environment variables for job 0.

Readers will note that selecting the node (708) having the JVM environment variables (652) that best matches the desired JVM environment variables (707) for the particular job (158) in the manner described above is for explanation only and not for limitation. In some other embodiments, the application manager may select the best match by taking into account the degree to which the JVM environment variables for a node match the desired JVM environment variables for a job in addition to the overall number of JVM environment variables for a node match the desired JVM environment variables. Still further, the application manager may select the best match by taking into weights for each type of JVM environment variable because some variables may be more important than others. Moreover, readers will note that selecting the node (708) having the JVM environment variables (652) that best matches the desired JVM environment variables (707) for the particular job (158) may be carried out in other ways as will occur to those of skill in the art.

After selecting (706) the node (708) on which to execute a particular job, the application manager may configure (710) that job (158) on the selected compute node (708) for execution according to the method of FIG. 7 by instructing a service node to write the particular job into computer memory of the selected node (708) and instructing the service node to in turn instruct the selected node (708) to begin executing the particular job (158) through that node's JVM. Although the configuring above occurs through a service node, readers will note that in some embodiments the application manager may itself provide the job to the selected node (708) and instruct the compute node to begin execution.

As described above, in some embodiments, an application manager tracks JVM environment variables for the JVMs installed on the compute nodes by receiving the JVM environment variables from each of the JVMs on the nodes. In such embodiments, the JVMs installed on the nodes record their individual JVM environment variables for reporting to the application manager. For further explanation, therefore, consider FIG. 8 that sets forth a flow chart illustrating a further exemplary method for executing a distributed JAVA application on a plurality of compute nodes according to embodiments of the present invention.

The method of FIG. 8 includes identifying (800), by each JVM supporting at least one job of the JAVA application on

the plurality of compute nodes, the JVM environment variables (652) for that JVM. Each JVM may identify (800) the JVM environment variables (652) for that JVM according to the method of FIG. 8 by logging in the JVM's environment variables (652) in a table, text file, JAVA object, or some other data structure as will occur to those of skill in the art.

The method of FIG. 8 includes providing (802), by each JVM supporting at least one job of the JAVA application on the plurality of compute nodes, the identified JVM environment variables (652) to the application manager (125). Each JVM may provide (802) the identified JVM environment variables (652) to the application manager (125) according to the method of FIG. 8 by transmitting the identified JVM environment variables (652) directly to the application manager (125) across a data communications network. In other embodiments, each JVM may provide (802) the identified JVM environment variables (652) to the application manager (125) according to the method of FIG. 8 by passing the identified JVM environment variables (652) to a network monitor installed on the node with the JVM. The network monitor may then pass the identified JVM environment variables (652) to the application manager (125) across a data communications network. In such a manner, the network monitor serves as a 'local client' for the application manager (125) on the node with the JVM. Upon receiving the JVM environment variables (652), the application manager (125) may then utilize the JVM environment variables (652) to configure jobs of a JAVA application on the plurality of nodes for execution according to embodiments of the present invention discussed above with reference to FIG. 7.

Exemplary embodiments of the present invention are described largely in the context of a fully functional computer system for executing a distributed JAVA application on a plurality of compute nodes. Readers of skill in the art will recognize, however, that the present invention also may be embodied in a computer program product disposed on computer readable media for use with any suitable data processing system. Such computer readable media may be transmission media or recordable media for machine-readable information, including magnetic media, optical media, or other suitable media. Examples of recordable media include magnetic disks in hard drives or diskettes, compact disks for optical drives, magnetic tape, and others as will occur to those of skill in the art. Examples of transmission media include telephone networks for voice communications and digital data communications networks such as, for example, Ethernet<sup>TM</sup> and networks that communicate with the Internet Protocol and the World Wide Web as well as wireless transmission media such as, for example, networks implemented according to the IEEE 802.11 family of specifications. Persons skilled in the art will immediately recognize that any computer system having suitable programming means will be capable of executing the steps of the method of the invention as embodied in a program product. Persons skilled in the art will recognize immediately that, although some of the exemplary embodiments described in this specification are oriented to software installed and executing on computer hardware, nevertheless, alternative embodiments implemented as firmware or as hardware are well within the scope of the present invention.

It will be understood from the foregoing description that modifications and changes may be made in various embodiments of the present invention without departing from its true spirit. The descriptions in this specification are for purposes of illustration only and are not to be construed in a limiting sense. The scope of the present invention is limited only by the language of the following claims.

What is claimed is:

1. A method of executing a distributed application on a plurality of compute nodes, the application comprising a plurality of jobs distributed among the plurality of compute nodes, the plurality of compute nodes connected together for data communications through a data communication network, each of the plurality of compute nodes having installed upon it a Virtual Machine ('VM') capable of supporting at least one job of the application, the method comprising:

tracking, by an application manager, VM environment variables for the VMs installed on the plurality of compute nodes; and

configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes.

2. The method of claim 1 wherein configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes further comprises performing for at least one job of the application:

selecting, for that job, one of the compute nodes on which to execute that job in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes and at least one desired VM environment variable for that job; and

configuring that job on the selected compute node for execution.

3. The method of claim 1 wherein the VM environment variables further comprise minimum and maximum heap sizes.

4. The method of claim 1 wherein the VM environment variables further comprise page size options.

5. The method of claim 1 wherein the VM environment variables further comprise debug trace status.

6. The method of claim 1 further comprising:

identifying, by each VM supporting at least one job of the application on the plurality of compute nodes, the VM environment variables for that VM; and

providing, by each VM supporting at least one job of the application on the plurality of compute nodes, the identified VM environment variables to the application manager.

7. The method of claim 1 wherein executing a distributed application on a plurality of compute nodes further comprises executing a distributed application on a parallel computer, the parallel computer comprising the plurality of compute nodes and a service compute node, the application manager installed upon the service compute node, the plurality of compute nodes connected for data communications through a plurality of data communications networks, at least one data communications network optimized for collective operations, and at least one other data communications network optimized for point to point operations.

8. A computing system capable of executing a distributed application on a plurality of compute nodes, the application comprising a plurality of jobs distributed among the plurality of compute nodes, the plurality of compute nodes connected together for data communications through a data communication network, each of the plurality of compute nodes having installed upon it a Virtual Machine ('VM') capable of supporting at least one job of the application, the computing system comprising one or more computer processors and computer memory operatively coupled to the computer processors, the computer memory having disposed within it computer program instructions capable of:

## 21

tracking, by an application manager, VM environment variables for the VMs installed on the plurality of compute nodes; and

configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes.

9. The computing system of claim 8 wherein configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes further comprises performing for at least one job of the application:

selecting, for that job, one of the compute nodes on which to execute that job in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes and at least one desired VM environment variable for that job; and

configuring that job on the selected compute node for execution.

10. The computing system of claim 8 wherein the VM environment variables further comprise minimum and maximum heap sizes.

11. The computing system of claim 8 wherein the VM environment variables further comprise page size options.

12. The computing system of claim 8 wherein the computer memory has disposed within it computer program instructions capable of:

identifying, by each VM supporting at least one job of the application on the plurality of compute nodes, the VM environment variables for that VM; and

providing, by each VM supporting at least one job of the application on the plurality of compute nodes, the identified VM environment variables to the application manager.

13. A computer program product for executing a distributed application on a plurality of compute nodes, the application comprising a plurality of jobs distributed among the plurality of compute nodes, the plurality of compute nodes connected together for data communications through a data communication network, each of the plurality of compute nodes having installed upon it a Virtual Machine ("VM") capable of supporting at least one job of the application, the

## 22

computer program product disposed upon a computer readable recordable medium, the computer program product comprising computer program instructions capable of:

tracking, by an application manager, VM environment variables for the VMs installed on the plurality of compute nodes; and

configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes.

14. The computer program product of claim 13 wherein configuring, by the application manager, the plurality of jobs for execution on the plurality of compute nodes in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes further comprises performing for at least one job of the application:

selecting, for that job, one of the compute nodes on which to execute that job in dependence upon the VM environment variables for the VMs installed on the plurality of compute nodes and at least one desired VM environment variable for that job; and

configuring that job on the selected compute node for execution.

15. The computer program product of claim 13 wherein the VM environment variables further comprise minimum and maximum heap sizes.

16. The computer program product of claim 13 wherein the VM environment variables further comprise page size options.

17. The computer program product of claim 13 wherein the VM environment variables further comprise debug trace status.

18. The computer program product of claim 13 further comprising computer program instructions capable of:

identifying, by each VM supporting at least one job of the application on the plurality of compute nodes, the VM environment variables for that VM; and

providing, by each VM supporting at least one job of the application on the plurality of compute nodes, the identified VM environment variables to the application manager.

\* \* \* \* \*